

A Region Based Easy Path Wavelet Transform For Sparse Image Representation

Renato Budinich

*Institut für Numerische und Angewandte Mathematik, Universität
Göttingen
Lotzestrasse 16-18, 37083 Göttingen, Deutschland
r.budinich@math.uni-goettingen.de*

February 16, 2017

Abstract

The *Easy Path Wavelet Transform* is an adaptive transform for bivariate functions (in particular natural images) which has been proposed in [13]. It provides a sparse representation by finding a path in the domain of the function leveraging the local correlations of the function values. It then applies a one dimensional wavelet transform to the obtained vector, decimates the points and iterates the procedure. The main drawback of such method is the need to store, for each level of the transform, the path which vectorizes the two dimensional data. Here we propose a variation on the method which consists of firstly applying a segmentation procedure to the function domain, partitioning it into regions where the variation in the function values is low; in a second step, inside each such region, a path is found in some deterministic way, i.e. not data-dependent. This circumvents the need to store the paths at each level, while still obtaining good quality lossy compression. This method is particularly well suited to encode a Region of Interest in the image with different quality than the rest of the image.

1 Introduction

In recent years there has been a big interest in using wavelet-like transforms for natural image compression and denoising. One wishes to find an appropriate transform such that the most important information of the image is concentrated in few coefficients; by then thresholding coefficients and keeping only the largest ones, one hopes to obtain a good quality approximation of the original image. However simply using a two-dimensional tensor wavelet transform doesn't yield good results, mainly because edges are poorly preserved. This is due to the

support of the basis elements not being adapted to the directional geometric properties present in natural images.

Therefore transforms such as curvelets [2, 3] and shearlets [6, 5] have been developed, which are highly redundant frames with strong anisotropic selectivity. However these are non-adaptive frames which loose their near optimal approximation properties if strong hypotheses are dropped on the edges in the image (namely piecewise C^2).

Instead of choosing an a priori frame to approximate the image, one can instead choose to adapt the frame elements to the particular image. Many different approaches to this concept have been studied in the recent years, such as bandlets [9], grouplets [11] and dictionary learning methods [1, 20]. For a review of adaptive image representations see [12].

In [13] a new adaptive transform for images termed *Easy Path Wavelet Transform* (EPWT) was introduced. In this method, a path is found among the points of the image that leverages local correlation in the gray values so as to produce a one dimensional signal with low entropy. At every level of the transform such a path is found and then a one dimensional wavelet transform is applied. The image quality obtained with this method is very good when compared to other methods; the main drawback comes from the need to store the paths for each level, which are needed during decoding. In [15] and [14] it was shown that, with a suitable choiche of the paths, the N -term approximation given by the EPWT is optimal for piecewise-Hölder functions. In [16] the EPWT was used as part of a hybrid method for Image Approximation while in [7] for denoising of scattered data.

Here we propose a variation on the original EPWT method, which we call the *Region Based Easy Path Wavelet Transform* (RBEPWT) method. The objective is to reduce the adaptivity storage cost by not requiring to store the paths like in the EPWT. In order to achieve this, a segmentation method is applied to the image in a first step, in order to partition the image into areas of low variation of the gray values. Then, for each region, a path is found in some kind of canonical manner: the path depends only on the geometrical shape of the region's border and not on the gray-values of the inner pixels. In this way in the final encoding one needs to store only the regions obtained from the initial segmentation step and the wavelet coefficients; the paths, being a deterministic function of the region's shape, can be recomputed on-the-fly during decoding, and thus need not be stored.

The quality of the lossy compression obtained from decoding a hard-thresholded set of coefficients heavily depends on the initial segmentation. In this regard we require a segmentation that finds regions where the local variance in the gray-values of the pixels is low - we do not care if the identified regions correspond to semantic areas in the image (which is the objective of many segmentation methods used in computer vision), but we wish instead to not have big jumps in the gray-values in the regions. We thus need a segmentation algorithm that essentially does a clustering: if we think of the image as a set of points in a three-dimensional space, where the first two coordinates are the indexes of the pixel and the third its gray-value, we wish for the segmentation to identify the

main clusters of such points.

In all our preliminary numerical tests we used the Felzenszwalb-Huttenlocher algorithm proposed in [4], because of its excellent quality to computation speed ratio. In the future we would like to further investigate this important first step in our method, and develop an ad hoc procedure to find regions that are already somehow optimal for the further steps of the RBEPWT transform. We will comment more on this in Section 2.2.3.

The outline of this paper is as follows. In Section 2 we introduce a common framework for the EPWT and RBEPWT. In Sections 2.1 and 2.2 we specify the particular details of the EPWT and RBEPWT respectively when viewed in this framework; in particular in Section 2.2.3 we comment on the desirable characteristics of a segmentation method for the RBEPWT. Finally in Section 3 we present the results of various numerical experiments.

2 Easy-path wavelet transform methods

Both the original EPWT method and the here proposed RBEPWT conform to a common general framework. Let $f : I \rightarrow \{0, \dots, 255\}$ be the input image, where $I = \{(i, j) \mid 0 \leq i, j \leq N - 1\}$ is the set of indexes and $N = 2^k$ for some $k \in \mathbb{N}$; assume I is ordered in some canonical manner, for example using the lexicographical ordering. Let $L \in \mathbb{N}$ be the number of levels of the transform; $2^L \leq N^2$ must hold. Choose a set of low-pass and high-pass 1-dimensional wavelet synthesis and reconstruction filters; following notation from [10] we will call them h, g, \bar{h}, \bar{g} .

Define $I^L := I$ and $f^L := f$. The encoding for the first level consists of:

1. Finding a *path* in I^L ; this can be thought equivalently as a function

$$p^L : \{0, 1, \dots, |I| - 1\} \rightarrow I^L \quad (2.1)$$

or as a permutation of the elements of I^L . How this path is found is the central point of such methods, and what differentiates the EPWT from the RBEPWT; we will comment more on this later. We can then define $\tilde{f}^L : I^L \rightarrow \{0, \dots, 255\}$ by $\tilde{f}^L := f^L \circ p^L$ as the 1-dimensional signal obtained from sampling the image along the path.

2. Apply one level of a periodic 1-dimensional discrete wavelet transform to \tilde{f}^L , obtaining approximation and detail coefficients a^L and d^L respectively:

$$\begin{aligned} a^L(k) &:= \tilde{f}^L \star \bar{h}(2k) \\ d^L(k) &:= \tilde{f}^L \star \bar{g}(2k), \quad k = 0, \dots, \frac{|I^L|}{2} - 1 \end{aligned} \quad (2.2)$$

where \star denotes the discrete convolution.

3. Define $I^{L-1} := \{p^L(2k) \text{ s.t. } k = 0, \dots, \frac{|I^L|}{2} - 1\}$; it is the subset of I^L obtained by taking only the coordinates corresponding to even indexes in

the path p^L , i.e. it is I^L decimated by a factor of 2, following the order induced by p^L . Define $f^{L-1} : I^{L-1} \rightarrow \{0, \dots, 255\}$ by $f^{L-1}(p^L(2k)) = a^L(k)$.

Now, I^{L-1} and f^{L-1} are a new set of indexes and a vector of values respectively, with half the points of I^L . One can thus iterate the 3 steps of encoding, at each level halving the number of points.

In summary the encoding steps of the transform consist of:

1. permuting the order of the points (i.e. find a path),
2. applying a discrete wavelet transform,
3. storing the detail vector and use the approximation vector as values for the new level, which is obtained by down-sampling the points along the found path.

Since we are interested in lossy compression, typically after the encoding a thresholding procedure will be used on the coefficients. Following the usual assumption that the coefficients with smallest absolute value in the encoded image are the least significant ones, we simply keep the n largest coefficients (in absolute value) and set the others to 0.

For decoding one needs the approximation vector for the lowest level, all the wavelet detail vectors and the permutations for each level. Then the decoding procedure consists simply of

1. applying the inverse wavelet transform for that level,
2. applying the inverse permutation for that level.

Note that, since we are only interested in the final result of the decoding (where the spatial disposition of the points is given by the canonical ordering chosen in advance), we do not need to know, for the intermediate levels, which value in the vector corresponds to which point in space. In other words, while during encoding we down-sample both the points in space and the vectors of values associated to them, in the decoding phase we can operate on the vectors only, upsampling and permuting them. The spatial information is needed only during encoding in order to find the path, i.e. the order according to which we vectorize the 2-dimensional data; during decoding this order (given by the inverse permutations) is already known.

2.1 EPWT

In the EPWT method the path at level l starts from a canonical point (for example $(0, 0)$) and at each step, among the closest points that are still available (i.e. aren't already part of the path), it makes a greedy choice: the point that

gives the least difference in absolute value is chosen

$$\begin{cases} p^l(0) & := (0, 0) \\ p^l(k+1) & \in \operatorname{argmin}_{(i,j) \in I^l \cap B_h^\circ(p^l(k))} |f^l(i, j) - f^l(p^l(k))| \end{cases} \quad (2.3)$$

where $\tilde{h} := \min \{h \in \mathbb{N}, B_h^\circ(p^l(k)) \not\subseteq \cup_{m=1}^{k-1} p^l(m), B_h^\circ(p^l(k)) \cap I^l \neq \emptyset\}$.

Here $B_h^\circ((i, j))$ denotes the punctured ball¹ centered in point (i, j) of radius h and \tilde{h} is the minimum required radius to obtain a neighborhood of point (i, j) that contains points in I^l that have not yet been taken as previous points of the path. If the argmin set contains more than one element, the choice can be done so as to minimize the direction change in the path.

In order to have access to the inverse permutations during encoding, the paths for each level have to be stored, alongside the wavelet detail and approximation coefficients. In [13] it was shown that, with the same number of coefficients, the EPWT greatly outperforms the classical 2-dimensional wavelet transform. However the storage cost is strongly affected by the need to store all the permutations; this is the point we wish to address here with the new RBEPWT method.

2.2 RBEPWT

In the region based EPWT, we first apply a segmentation method to the image, in order to obtain a subdivision of the original image into *regions* (subsets of points). The rationale is that if we have regions where the variation in gray-values is small, then it is not so important which path is taken inside that region. We can then have a canonical path-finding procedure which does not depend on the gray-values. In this way the final encoding consists of the wavelet detail, the approximation coefficients and the segmentation. We need not to store all the paths, since these depend only on the segmentation and not on the pixel grayvalues and thus they can be recomputed on-the-fly during decoding.

We will comment in Section 2.2.3 on the properties the segmentation method should have. For now we suppose that the segmentation step identifies regions $R_0, R_1, \dots, R_{r-1} \in \mathcal{P}(I)$ where I is the set of indexes in the image and $\mathcal{P}(I)$ denotes the set of sets of I . The regions are given in the form of a *label image* $\Lambda \in \{0, 1, \dots, r-1\}^{N \times N}$, obtained by filling region R_h with the value h , i.e. such that $(i, j) \in R_h \Leftrightarrow \Lambda_{i,j} = h$.

Suppose now we have a function Π that associates to any set of points a Hamiltonian path (in the complete graph generated by these points). In other words, for any region R we wish for $\Pi(R)$ to be a bijection from $\{0, 1, \dots, |R| - 1\}$ to R . Later we will present two examples of such functions, the easy-path and the grad-path.

Call $R_k^L := R_k$ for all $k = 0, \dots, r-1$ and define the region collection at the highest level as $\mathcal{R}^L := \{R_0^L, R_1^L, \dots, R_{r-1}^L\}$. Define for each region k the path

¹we use either the max distance, defined by $d((i, j), (k, l)) = \max \{|i - k|, |j - l|\}$, or the usual Euclidean distance

in it as $\rho_k^L := \Pi(R_k^L)$. By gluing all these paths together we obtain

$$p^L := \rho_0^L \cup \rho_1^L \cup \dots \cup \rho_{r-1}^L,$$

which is a permutation of the elements of I^L . At this point a global path is defined, and we can proceed just as described in Section 2 for the general framework: define $\tilde{f}^L := f^L \circ p^L$, compute a^L and d^L through discrete convolution with the two filters, define I^{L-1} and f^{L-1} . Additionally now we have to define the new region collection $\mathcal{R}^{L-1} := \{R_0^{L-1}, R_1^{L-1}, \dots, R_{r-1}^{L-1}\}$, where $R_k^{L-1} = R_k^L \cap I^{L-1}$. An equivalent definition would be:

$$\begin{aligned} R_0^{L-1} &= \rho_0^L(\text{even numbers}) \\ R_k^{L-1} &= \begin{cases} \rho_k^L(\text{even numbers}) & \text{if } |R_{k-1}^{L-1}| \text{ is even } \vee R_{k-1}^{L-1} = \rho_{k-1}^L(\text{odd numbers}) \\ \rho_k^L(\text{odd numbers}) & \text{otherwise,} \end{cases} \end{aligned} \quad (2.4)$$

where here \vee denotes the exclusive OR, i.e. it evaluates to true when only one of the two conditions is true. This procedure can be repeated analogously for each level. As already mentioned, the final encoding of the image will consist of the segmentation information, all the wavelet detail coefficients and the approximation coefficients for the lowest level.

For decoding, since the permutations here aren't stored, we first need to recompute them; to do so we simply need to apply the whole encoding procedure ignoring the pixel values but not the segmentation information, which has been stored and is thus available.

By what has been said it is clear that for our method, the path finding procedure Π must have the following characteristics:

1. It must not depend on the points in the region being uniform: from level $L-1$ onward in fact the points in each region will usually not be uniformly distributed.
2. It must be completely deterministic: this is needed in order to obtain the same paths during encoding and decoding.
3. It must not depend on the gray-values of the image: these aren't available during decoding. An exception can be made, as in the grad-path, if one is willing to store additional information in the final encoding.

2.2.1 Easy-Path

Algorithm 1 shows the easy-path procedure to find a path π in a region R : starting from some point (which for example can be chosen using the lexicographical ordering), the algorithm tries always to select the closest available neighbour. If there are more points equally close, it selects the one that would make for the straightest path, the rationale being that a more regular path will lead to a smoother signal (see for example [7] for a proof that, when f is sufficiently

Algorithm 1 Easy-Path algorithm

Input: region R

Output: path π

```

    choose starting point  $p \in R$ 
     $Q = R \setminus p$ 
     $v = (1, 0)$ 
    while  $Q \neq \emptyset$  do
5:    $\bar{h} = \min \{h \in \mathbb{N}, B_h(p) \cap Q \neq \emptyset\}$ 
       $C = \operatorname{argmin}_{\kappa \in B_{\bar{h}}(p) \cap Q} \|p - \kappa\|$ 
      if  $|C| \geq 2$  then
         $C = \operatorname{argmax}_{\phi \in C} \langle \phi - p, v \rangle$ 
        if  $|C| \geq 2$  then
10:      rotate  $v$  by  $\pi/2$ 
          goto 8
        end if
      end if
      pick  $\psi \in C$ 
15:   append  $\psi$  to  $\pi$ 
       $v = \psi - p$ 
       $p = \psi$ 
      remove  $\psi$  from  $Q$ 
    end while
20: return  $\pi$ 

```

$\triangleright p$ is the current point

$\triangleright Q$ is the set of available points

$\triangleright v$ is the preferred direction

\triangleright there is no choice to be done here: $C = \{\psi\}$

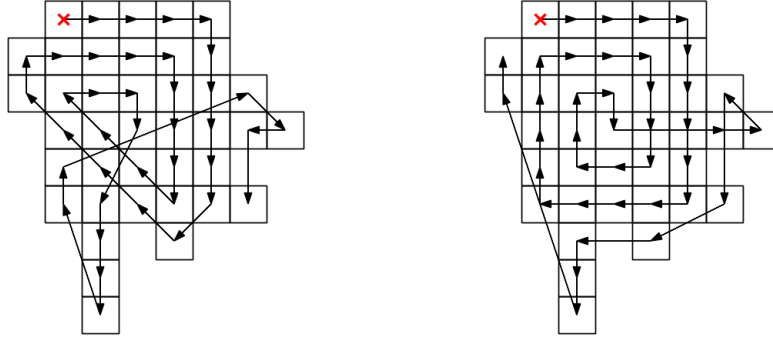


Figure 1: Path found by the easy-path procedure with the max (left) and Euclidean (right) distance

smooth, a straighter path gives smaller wavelet coefficients). This is done by computing the scalar product of the increment vector with a preferred direction vector, which at every iteration is updated to be the last increment in the path. If there are 2 possible points with the same minimum angle to the preceeding part of the path, then the preferred direction is rotated by $\pi/2$, making then only one of the two points preferable.

See Figure 1 for an example of a path determined by easy-path. In our tests using the Euclidean distance gave better compression performance, so we used that in all numerical results presented here.

2.2.2 Grad-path

Another path-finding procedure is the grad-path, shown in Algorithm 2. It requires previous computation of the average discretized gradient for each region; these vectors have to be stored, contributing to the storage cost of the final encoding. The procedure is very similar to the easy-path: the closest point is always preferred. However the preferred direction is always perpendicular to the average gradient, at each iteration the sign being updated so as to obtain the most regular path. Furthermore taking the absolute value when computing the scalar product (see line 8) means that we always prefer a path that remains as much as possible perpendicular to the average gradient, even if it means a sharper change in the path's direction. Only in case of equally distant points forming equal angles the preferred direction is temporarily rotated.

See Figure 2 for an example of the path generated by the gradpath algorithm. The points forming the region are the same in the two images, but the gray-values generate two different average gradients.

In our numerical tests (see Section 3) the grad-path performed only marginally

Algorithm 2 Grad-path algorithm

Input: region R , average gradient g

Output: path π

```

    choose starting point  $p \in R$  ▷  $p$  is the current point
     $Q = R \setminus p$  ▷  $Q$  is the set of available points
     $v = \text{rotate } g \text{ by } \pi/2$  ▷  $v$  is perpendicular to the average gradient
    while  $Q \neq \emptyset$  do
5:    $\bar{h} = \min \{h \in \mathbb{N}, B_h(p) \cap Q \neq \emptyset\}$ 
       $C = \underset{\kappa \in B_{\bar{h}}(p) \cap Q}{\operatorname{argmin}} \|p - \kappa\|$ 
      if  $|C| \geq 2$  then
         $C = \underset{\phi \in C}{\operatorname{argmax}} |\langle \phi - p, v \rangle|$ 
        if  $|C| \geq 2$  then
10:   rotate  $v$  by  $\pi/2$ 
        goto 8
      end if
    end if
    pick  $\psi \in C$  ▷ there is no choice to be done here:  $C = \{\psi\}$ 
15:  append  $\psi$  to  $\pi$ 
       $w = \text{rotate } g \text{ by } \pi/2$ 
       $v = \underset{\gamma = -w, w}{\operatorname{argmax}} \langle \psi - p, \gamma \rangle$ 
       $p = \psi$ 
      remove  $\psi$  from  $Q$ 
20: end while
    return  $\pi$ 

```

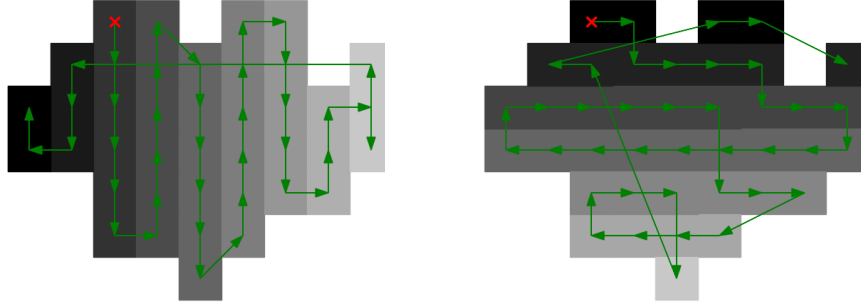


Figure 2: Gradpath on same region with different gray-values

better than the easy-path - this is without taking into account the additional storage cost.

2.2.3 Notes on the segmentation procedure

In order to avoid the need to store the permutation at each level as happens in the EPWT, in the proposed RBEPWT we are finding paths that do not depend on the data, at least not directly². In fact, the segmentation uses the gray-value information, but once the regions are determined the paths along which we sample are completely determined by the regions and are agnostic of the gray-value data. It is thus clear that the segmentation procedure is of primary importance for our method.

Given a segmentation $\mathcal{R} = \{R_0, \dots, R_{r-1}\}$ of the index set I , we can define the *perimeter* of \mathcal{R} as

$$P(\mathcal{R}) := \left| \left\{ ((i, j), (k, l)) \text{ s.t. } \|(i, j) - (k, l)\| = 1, (i, j) \in R_a, \right. \right. \\ \left. \left. (k, l) \in R_b \text{ for some } a \neq b \right\} \right|,$$

where we consider the Euclidean norm. In other words the perimeter is the number of edges on the border between the regions of the segmentation - here we are thinking of an edge at the region's boundary as a pair of pixels on opposite sides of it, which have distance 1 in the Euclidean norm. If we assume the regions are all path connected when considering only paths that move horizontally or vertically in the image (i.e. not considering paths with segments of the type $(i, j), (i + 1, j + 1)$), then these bordering edges completely define

²In this section we are assuming that the easy-path procedure from Section 2.2.1 is being used

the segmentation. In any case, the perimeter gives a first relative estimate of the storage cost of the segmentation: in fact, though we haven't yet explicitly determined an efficient encoding of the perimeter edges, we assume that it will somehow encode the direction changes of the regions' border, such that e.g. a long straight border will be encoded as a sequence of 0 and thus will be efficiently compressible.

Thus, having a small perimeter and regular borders are desirable characteristics of a segmentation for our method. A more fundamental goal is to produce regions such that, when its points are reordered according to the easy-path procedure, the gray-values considered in this order give a sequence with few jumps. A potentially tricky situation would be the image depicted in Figure 3 (a): segmentations (b) and especially (c) are here clearly preferable to (d), because the easy-path procedure applied to the latter one would pass many times through the central vertical line of the image and produce a signal with many jumps.

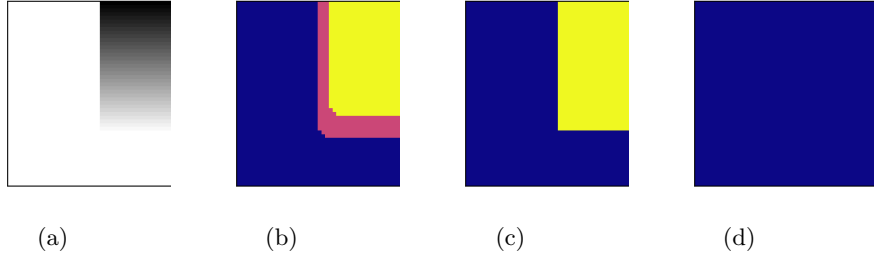


Figure 3: (a) Original image, (b) Felzenszwalb-Huttenlocher segmentation for some scale parameter k_0 and smoothing parameter $\sigma > 0$, (c) $\sigma = 0$, and finally (d) for $\sigma = 0$ and scale parameter $k_1 \gg k_0$.

In our tests we always applied Felzenszwalb-Huttenlocher's segmentation proposed in [4]³, because of its very good performance to computation speed ratio. This method considers the graph $G = (V, E)$, where the set of vertices are the pixels of the images and the edges connect each pixel to the 8 pixels forming its square neighborhood. The weights $w(e)$ of the edges e are given by the difference in absolute value of the gray-values. Given any region (set of vertices) C , its *internal difference* is defined as

$$C \subseteq V, \text{Int}(C) := \max_{e \in \text{MST}(C, E)} w(e), \quad (2.5)$$

where $\text{MST}(C, E)$ is the minimum spanning tree of the subgraph with vertices C and all edges in E that are between vertices in C , i.e. $(C, E \cap C \times C)$. This

³more precisely its implementation in the **skimage** python library

quantity measures the inner variance in gray-value, and tells us that there is a path between any two pixels in C using only edges of weight smaller than $\text{Int}(C)$. Furthermore, given two regions C_1 and C_2 , the authors in [4] define the *difference* between the two regions as:

$$C_1, C_2 \subseteq V, \text{Diff}(C_1, C_2) := \min_{(v_i, v_j) \in C_1 \times C_2 \cap E} w((v_i, v_j)) . \quad (2.6)$$

Finally, given two regions C_1 and C_2 , they define a predicate, which evaluates to **True** if there is evidence of a boundary between the two regions and to **False** otherwise:

$$D(C_1, C_2) := \begin{cases} \text{True} & \text{if } \text{Diff}(C_1, C_2) > \min \{I_k(C_1), I_k(C_2)\} \\ \text{False} & \text{otherwise} \end{cases} , \quad (2.7)$$

where

$$I_k(C) := \text{Int}(C) + \frac{k}{|C|} \quad (2.8)$$

is the *scaled internal difference* of a region. The scaling parameter k is introduced to give a preference to regions of cardinality not too small with respect to k . Note that instead of the term $\frac{k}{|C|}$ any positive function $\tau(C)$ could be used: the segmentation will then give a preference to regions C for which $\tau(C)$ is big. In the future we wish to experiment with this function in order to optimize the regions for our method. One possible example would be to define $\tau(C) = \frac{\text{area}(C)}{\text{perimeter}(C)}$, in order to obtain a segmentation with a smaller global perimeter and thus smaller storage cost.

The algorithm then starts with each pixel defined as its own region and then, considering each edge in increasing weight, evaluates the predicate for the two regions connected by the edge. It is proven to produce a segmentation which is neither too coarse nor too fine; see [4] for details. Prior to this procedure a Gaussian filter with variance σ is applied to smooth the image.

When comparing this segmentation method against the requirements of a good segmentation for our method, as described in the first part of this Section, we must make two remarks. Firstly, the regions produced are such that if we use the MST of that region to move between points, we will use only edges of low weight. But when we apply our method, we will be moving along the path generated by the easy-path procedure, and thus we will not have any guarantee of not obtaining a signal with jumps. Secondly, since in (2.6) the minimum edge weight is considered, the method can behave badly if there is a portion of the image similar to Figure 3a of size much smaller than the scale parameter k . This could be avoided by substituting the minimum edge weight in (2.6) with the median edge weight or another quantile, but this is shown in [4] to make the segmentation problem NP-hard.

2.3 Thresholding for a Region of Interest

Though we have supposed the index set to be that of a square image, i.e. $I = \{(i, j) \mid 0 \leq i, j \leq N - 1\}$, there is nothing impeding us to relax this con-

dition: the RBEPWT method can be applied without modification to images with arbitrary boundary shape. In particular, it can be used to encode only a region of interest (ROI) of a whole image - see Figure 4. One could then

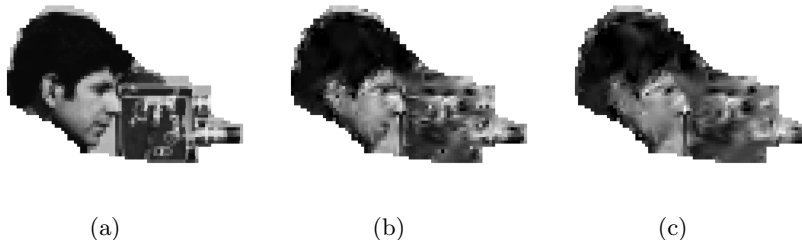


Figure 4: Encoding of a ROI for the cameraman image using the CDF 9/7 wavelet and 11 levels. (a): full quality with 2101 coefficients, (b): thresholded to 210 coefficients and (c): thresholded to 100 coefficients.

think of encoding with different qualities the ROI and the rest of the image, simply by viewing them as two separate images with non-regular boundaries. To this end however we can do something a little bit more clever, leveraging the fact that many coefficients of the full encoding of the image, especially at low levels, contain highly non-local information, shared by both the ROI and its complementary. This allows us fine control over the quality of the encoding for a region of arbitrary shape (and its complementary), something that is not possible with the classical 2-dimensional wavelet transform.

For simplicity we will restrict ourselves to the case of Haar wavelets and we will suppose that the region of interest R is one of the regions found in the segmentation step; we will also suppose the image to be of size $2^N \times 2^N$ and to encode using the maximum number of levels $L = 2 \log_2 2^N = 2N$. The key observation is that each value of f^{L-1} will be determined only by two values of f^L ; this follows from (2.2) and the fact that the low and high pass Haar filters have support length 2. In order to determine all coefficients in the encoding necessary to reconstruct R , it is useful to organize the coefficients in a graph structure, similarly as to how is done in [18], where a generalization of EPWT-like methods is proposed. There the authors define, for every level l of the transform, a node in the graph for each element in I^l ; edges are defined between a node from level l (i.e. a node associated to a point in I^l) to another at level $l - 1$ if the value of the latter is associated to a nonzero value of the filter in the sum defining the former. Thus this graph depends on the filter, and specifically only on the lengths of its support; in the Haar case, when the maximum number of levels is used, the graph is actually a full complete binary tree.

In order to visualize the permutation, we add to this construction for each level a set of nodes, again each representing a point in that level of the transform. We thus have for each point in I^l two nodes in the graph, and we can then view the permutation found at level l by the path-finding procedure of the RBEPWT

as a bipartite subgraph. See Figure 5 for an example.

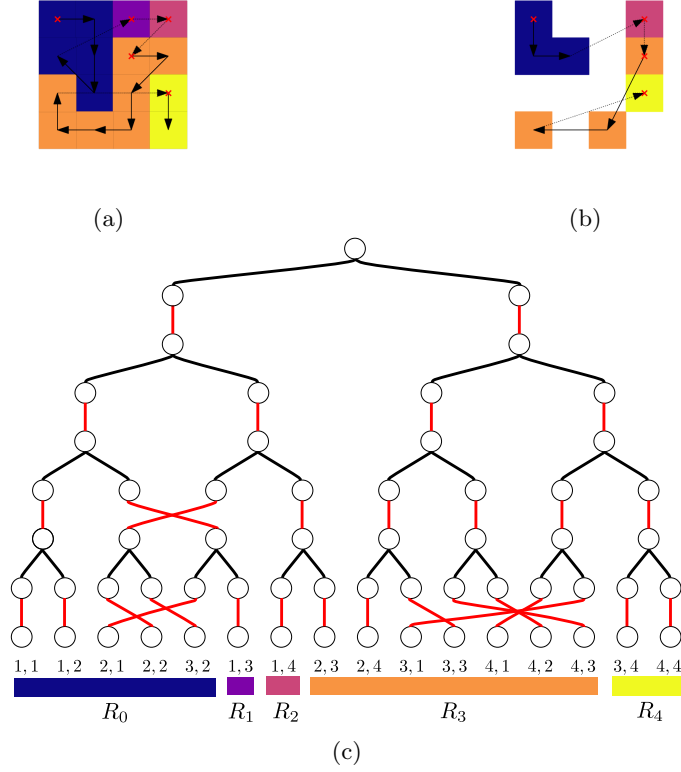


Figure 5: (a): A segmented 4×4 image with the path found by the easy path procedure at the fourth level, (b): the regions and path at the third level, (c): the coefficient tree associated to the RBEPWT with Haar wavelets and 4 levels. Here the implied order of the pixels is given by row-stacking (i.e. left to right and top to bottom). The leafs are denoted by their coordinates in the image and are grouped by region. The red edges represent the permutation given by the path finding procedure, while the black edges represent the convolution with the low and high pass filters.

Now, the points in R correspond to a set of leaf nodes in the tree; from the definition of the tree, it follows that only the RBEPWT coefficients corresponding to nodes ancestors of these leaf nodes are influenced by the values of the points in R . The farther up the tree a node is, the larger is its zone of influence in the image - the root node is affected by all points in the image, but closer to the leafs the coefficients contain more localized information. For example in Figure 5 5 coefficients (4 detail and one approximation) are needed to perfectly reconstruct R_4 , and of these only one (the root node) is common to R_0 . If we wish to encode only region R then it suffices to follow the edges

of the tree up until the root node, starting from the leafs associated to R and preserving the coefficients encountered in this visit while setting to 0 all the others. In this way all the information necessary to perfectly reconstruct R will be preserved alongside some information of the rest of the image; (see Figure 10b). To encode with different qualities R and the rest of the image, we must divide the coefficients into two sets: those that are ancestors to points in R in the tree and those that are not. Then we threshold the coefficients, preserving a certain percentage of those in the first set and another percentage of those in the second but not in the first, always giving precedence to coefficients larger in absolute value (see Figure 10c).

The general non-Haar case is more complicated: one point in I^l will influence $\lfloor s/2 \rfloor$ coefficients at level $l - 1$, where s is the length of the filter and is usually greater than 2. This means the graph will not be a tree anymore, since there will be leafs with more than one parent; furthermore in the non-orthogonal case the reconstruction filters may have different lengths than the synthesis, thus making it not obvious from the graph representation which coefficients are needed for perfect reconstruction of the ROI.

Finally, this idea could be applied to the standard EPWT and tensor wavelet transform as well, however both cases would require more coefficients to encode a ROI compared to the RBEPWT case. For the EPWT, the permutation at each level is searched globally on the whole set of points and not for each region separately, thus the coefficients could potentially get mixed more, with points that are nearby in the image sharing less common coefficients in the graph. For the tensor wavelet transform, the number of coefficients at level $l - 1$ influenced by a single coefficient at level l is of the order of $(s/2)^2$ ⁴, due to the nature of the transform that uses two-dimensional convolution.

3 Numerical Experiments

We present here some numerical results, which were all obtained using the python code available at the software page of <http://na.math.uni-goettingen.de/>⁵. We used the biorthogonal Cohen-Daubechies-Feauveau wavelets (known also as CDF 9/7) because, analogously to other natural images compression methods, they showed very good performance in our method compared to other wavelets. In Figures 6-8 our method is compared to the classical two-dimensional tensor wavelet transform and the EPWT. It is interesting to note how thresholding the coefficients introduces different types of distortion in the different transforms: the 2D wavelet transform blurs the image, while the EPWT and RBEPWT methods preserve edges much better, while introducing a higher frequency noise. In particular, as should be expected, the RBEPWT preserves the segmentation information: for example in the cameraman image, the skyscraper to the right of the camera tripod can be seen only in the RBEPWT images, because it was identified as a separate region in the segmentation step. In the

⁴supposing the low and high pass filter both have length s

⁵also available at <https://github.com/nareto/rbepwt>

house image the texture of the bricks is completely gone, but the borders are still perfectly visible - thus our method is particularly well-suited for cartoon-like images.

In order to quantify the quality of our method, we computed the peak signal to noise ratio (PSNR)

$$\text{PSNR}(f, \tilde{f}) := 20 \log_2 \frac{255}{\|f - \tilde{f}\|_2},$$

where here f is the matrix corresponding to the original image and \tilde{f} to the reconstruction using only the most significant RBEPWT coefficients. However, in recent years there has been a lot of research to find a better quality measurement index (for example see [21] and [22]). Here by better we mean that it correlates better to scores given by humans to different degraded versions of the same image, such as those presented in [17] or [8]. In [19], the authors show that on average their proposed HaarPSI index has the best correlation to such databases. Thus, in Figure 9, we show the PSNR (for comparison with other methods, since despite its shortcomings it is still widely used) and the HaarPSI values for the cameraman, peppers and house image. These plots show that our method, compared to EPWT, performs very good on the house image, which is almost cartoon-like, and a little worse on the cameraman and especially peppers image.

It should be noted here that a better comparison would be obtained by comparing the quality of the thresholded images with different transforms having the same storage cost (and not the same number of coefficients). However at this point it is unclear to us how the segmentation can be efficiently encoded.

In Figure 10 we show an example of the ROI based thresholding described in section 2.3.

Finally, we observe that by setting one coefficient to 1 and the rest to 0 and then decoding, it is possible to view the elements of the adaptive basis produced by our method. In Figure 11 some of such elements are shown for the cameraman image, both for the RBEPWT and the EPWT. As should be expected, the basis for our method clearly preserves the segmentation structure, while the basis for the EPWT appears to be totally unstructured.

4 Final Remarks

We introduced a new method for compression of natural images, which applies an image segmentation method to the image before applying an EPWT-like method. We defined two region path-finding procedures, which are deterministic and do not depend on the single gray values of the pixels in the region. We used the defined transform to compute a lossy compressed version of the original image, obtained by setting to 0 all but the greatest coefficients and applying

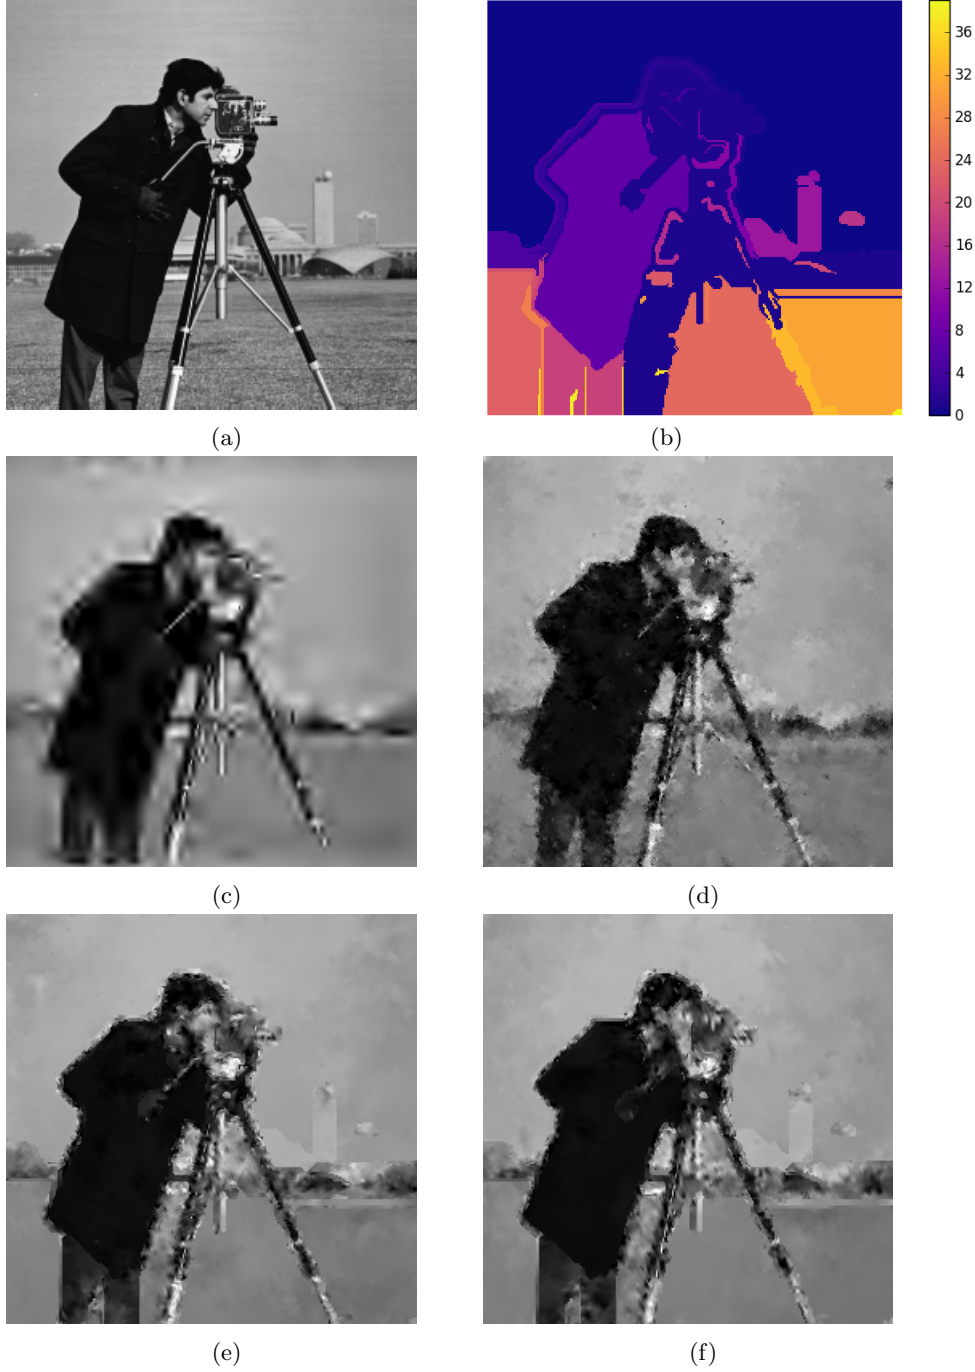


Figure 6: (a): Original cameraman image. (b): Segmentation obtained from the Felzenszwalb-Huttenlocher algorithm with scale parameter 200, $\sigma = 2$ and minimum scale 10. (c)-(f): Image compressed using 512 coefficients and the CDF 7/9 filter: (c) the classical 2D tensor wavelet transform, (d) EPWT, (e) RBEPWT with easy-path and (f) RBEPWT with grad-path transforms.

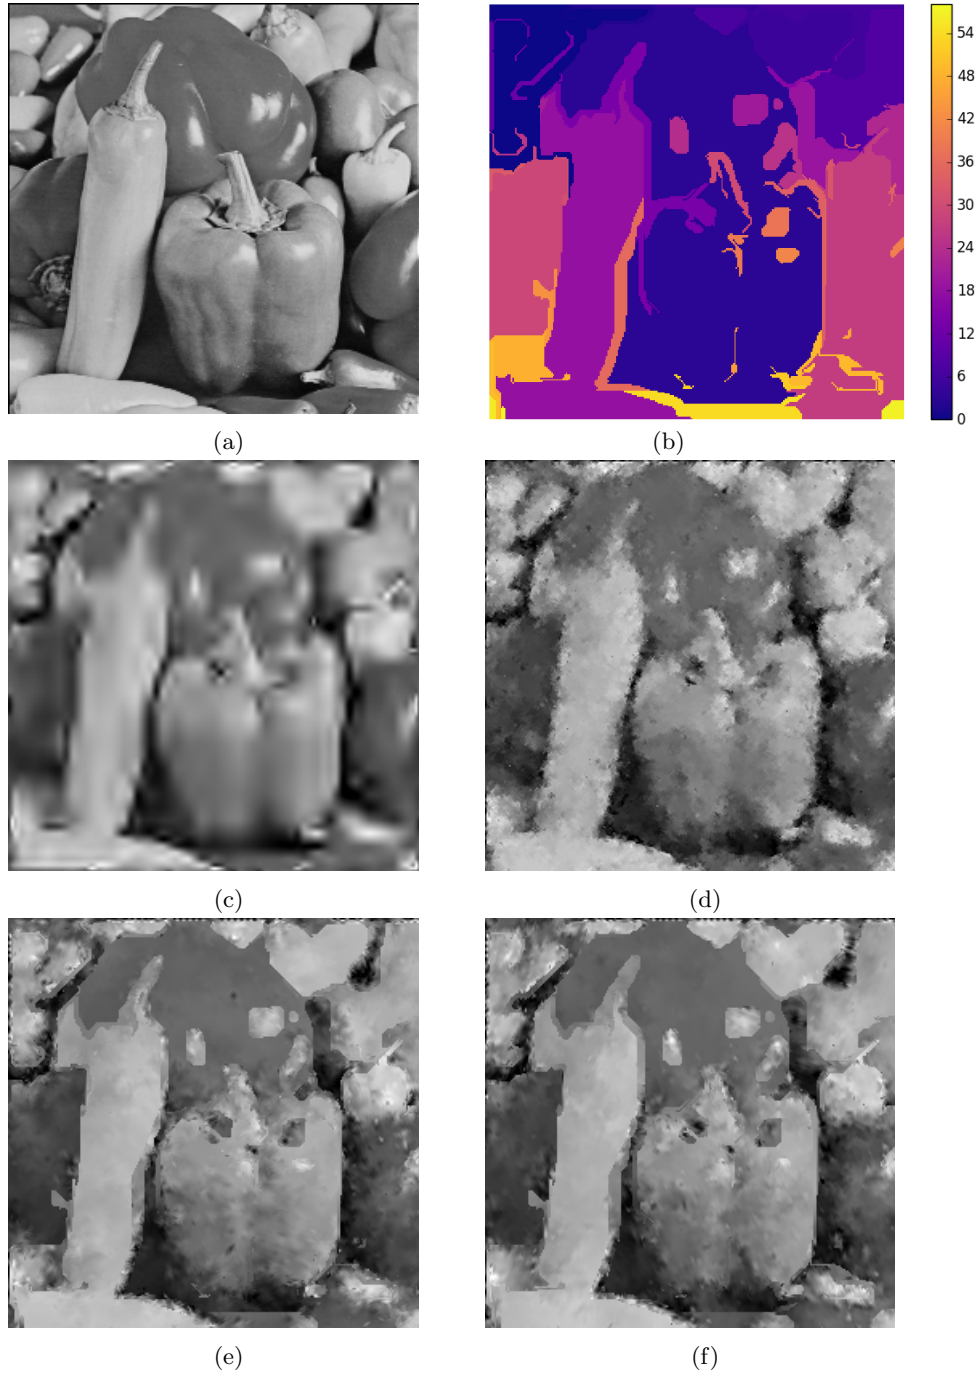


Figure 7: (a): Original peppers image. (b): Segmentation obtained from the Felzenszwalb-Huttenlocher algorithm with scale parameter 200, $\sigma = 2$ and minimum scale 10. (c)-(f): Image compressed using 512 coefficients and the CDF 7/9 filter: (c) the classical 2D tensor wavelet transform, (d) EPWT, (e) RBEPWT with easy-path and (f) RBEPWT with grad-path transforms.

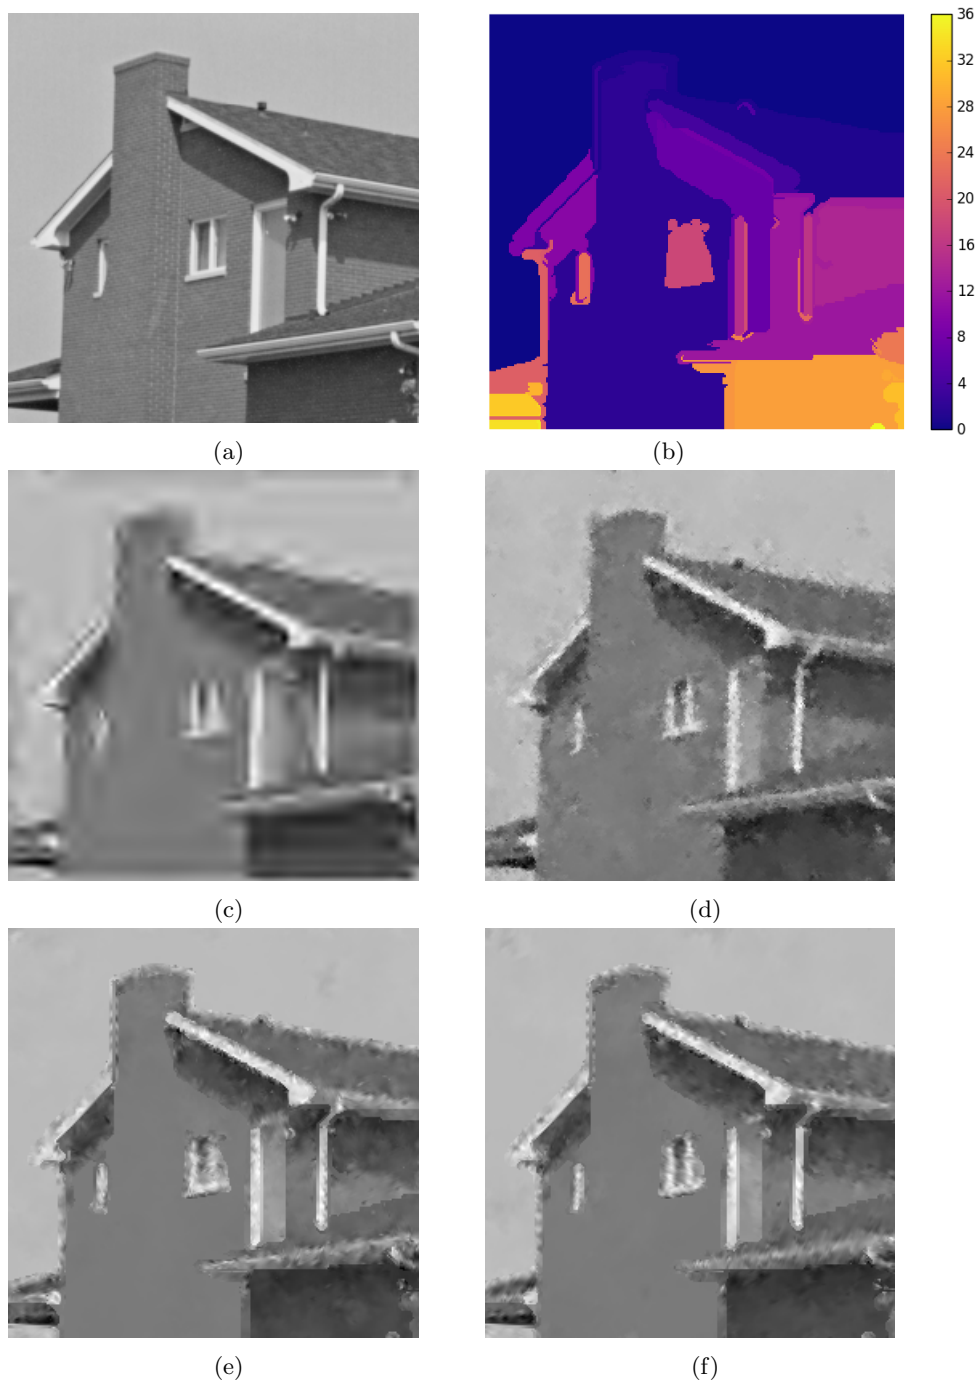
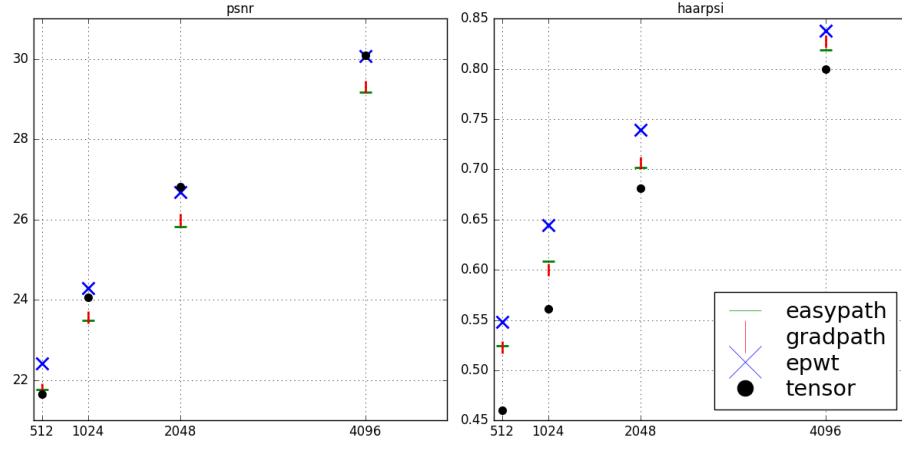
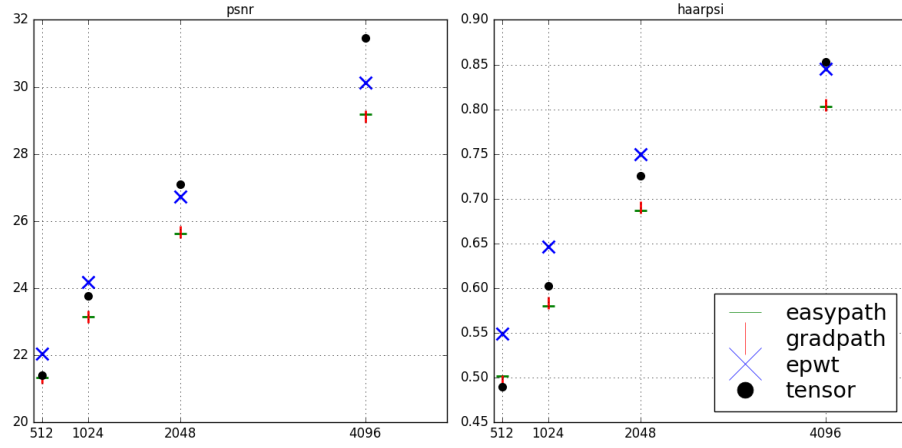


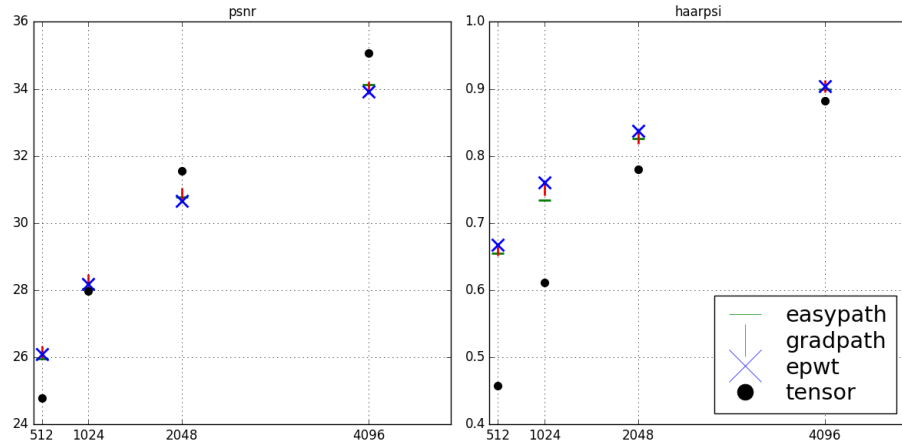
Figure 8: (a): Original house image. (b): Segmentation obtained from the Felzenszwalb-Huttenlocher algorithm with scale parameter 200, $\sigma = 2$ and minimum scale 10. (c)-(f): Image compressed using 512 coefficients and the CDF 7/9 filter: (c) the classical 2D tensor wavelet transform, (d) EPWT, (e) RBEPWT with easy-path and (f) RBEPWT with grad-path transforms.



(a)



(b)



(c)

Figure 9: PSNR and HaarPSI values for the (a) cameraman, (b) peppers and (c) house image, with 512, 1024, 2048 and 4096 coefficients

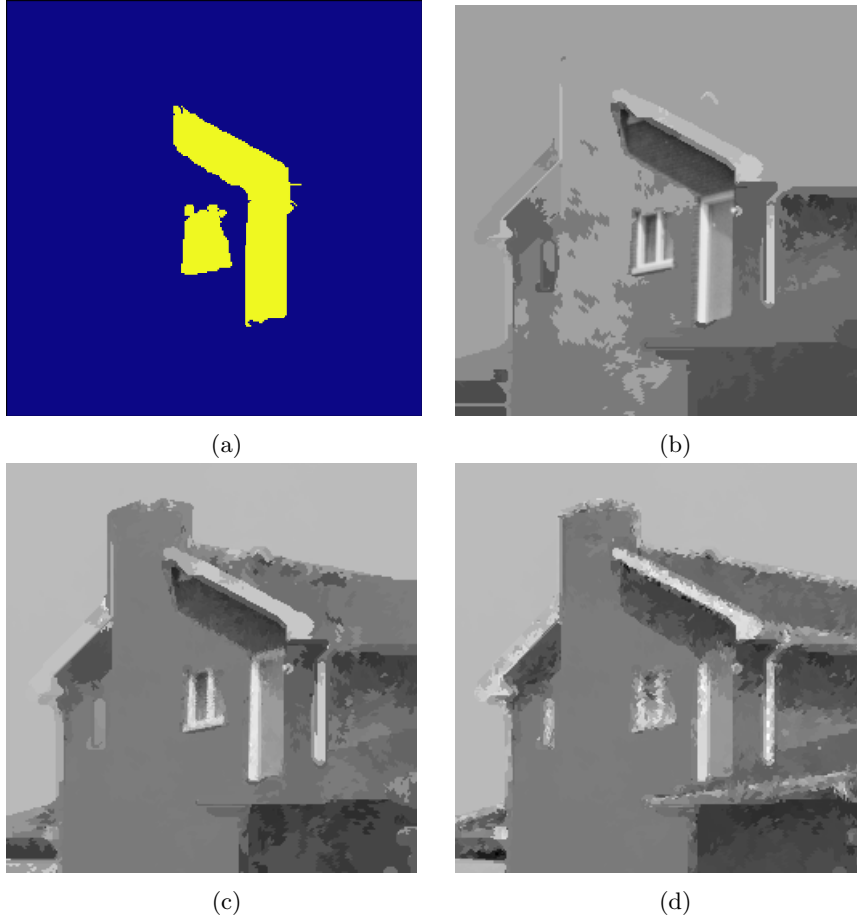
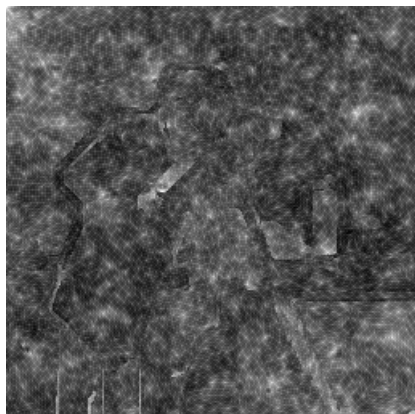


Figure 10: (a): the selected ROI for the house image, roughly corresponding to the window, the light part of the wall and the area under the roof; (b): decoded image where all and only the coefficients ancestors to the ROI were preserved (5042 non-zero coefficients); (c): decoded image where 10% of coefficients ancestors to the ROI and 0.1% not ancestors were preserved, for a total of 551 non-zero coefficients; (d): decoded image with standard thresholding with 551 non-zero coefficients (inserted here for comparison with (c))

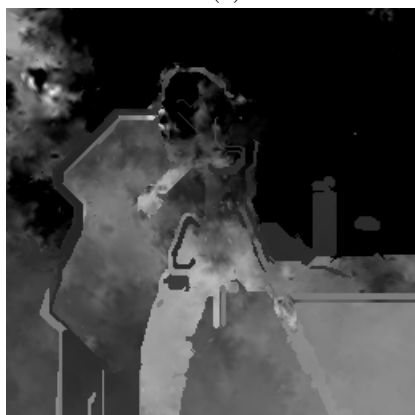
the decoding procedure. We commented on which are desirable characteristics for the segmentation step of our method, namely producing regions that enable paths generated by the easy-path procedure with few jumps and with small and regular perimeters. Finally we showed how the RBEPWT can be used to encode with different quality a specific region of interest from the rest of the image. In the future we wish to study more closely the segmentation step and possibly develop an ad-hoc procedure for our method.



(a)



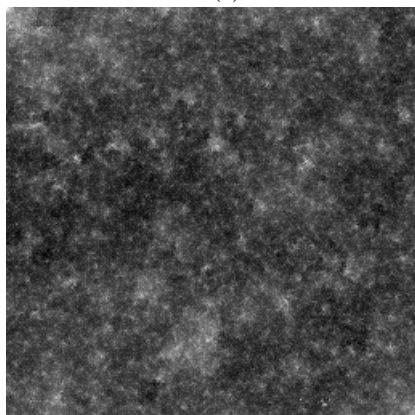
(b)



(c)



(d)



(e)



(f)

Figure 11: Basis elements for the cameraman image, encoded with 16 levels. (a)-(b): approximation and detail coefficients at the first level of the RBEPWT; (c)-(d) the two detail coefficients at the second level of the RBEPWT; (e)-(f) approximation and detail coefficients at the first level of the EPWT.

Acknowledgments

The author gratefully acknowledges the funding of this work by the DFG in the framework of the GRK 2088 and heartfully thanks Gerlind Plonka whose guidance and insight have been fundamental for the completion of this work.

References

- [1] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing*, 54(11):4311–4322, 2006.
- [2] Emmanuel Candes, Laurent Demanet, David Donoho, and Lexing Ying. Fast discrete curvelet transforms. *Multiscale Modeling & Simulation*, 5(3):861–899, 2006.
- [3] Emmanuel J Candès and David L Donoho. New tight frames of curvelets and optimal representations of objects with piecewise c^2 singularities. *Communications on pure and applied mathematics*, 57(2):219–266, 2004.
- [4] Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Efficient graph-based image segmentation. *Int. J. Comput. Vision*, 59(2):167–181, September 2004.
- [5] Kanghui Guo and Demetrio Labate. Optimally sparse multidimensional representation using shearlets. *SIAM journal on mathematical analysis*, 39(1):298–318, 2007.
- [6] Kanghui Guo, Demetrio Labate, Wang-Q Lim, Guido Weiss, and Edward Wilson. Wavelets with composite dilations. *Electronic research announcements of the American Mathematical Society*, 10(9):78–87, 2004.
- [7] Dennis Heinen and Gerlind Plonka. Wavelet shrinkage on paths for denoising of scattered data. *Results in Mathematics*, 62(3):337–354, 2012.
- [8] Eric C. Larson and Damon M. Chandler. Most apparent distortion: full-reference image quality assessment and the role of strategy. *J. Electronic Imaging*, 19(1):011006, 2010.
- [9] Erwan Le Pennec and Stephane Mallat. Bandelet image approximation and compression. *Multiscale Modeling & Simulation*, 4(3):992–1039, 2005.
- [10] S. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way*. Elsevier Science, 2008.
- [11] Stéphane Mallat. Geometrical grouplets. *Applied and Computational Harmonic Analysis*, 26(2):161–180, 2009.
- [12] Gabriel Peyré. A review of adaptive image representations. *IEEE Journal of Selected Topics in Signal Processing*, 5(5):896–911, 2011.

- [13] Gerlind Plonka. The easy path wavelet transform: A new adaptive wavelet transform for sparse representation of two-dimensional data. *Multiscale Modeling & Simulation*, 7(3):1474–1496, 2008.
- [14] Gerlind Plonka, Armin Iske, and Stefanie Tenorth. Optimal representation of piecewise hölder smooth bivariate functions by the easy path wavelet transform. *Journal of Approximation Theory*, 176:42–67, 2013.
- [15] Gerlind Plonka, Stefanie Tenorth, and Armin Iske. Optimally sparse image representation by the easy path wavelet transform. *International Journal of Wavelets, Multiresolution and Information Processing*, 10(01):1250007, 2012.
- [16] Gerlind Plonka, Stefanie Tenorth, and Daniela Rosca. A new hybrid method for image approximation using the easy path wavelet transform. *IEEE Transactions on Image Processing*, 20(2):372–381, 2011.
- [17] Nikolay Ponomarenko, Lina Jin, Oleg Ieremeiev, Vladimir Lukin, Karen Egiazarian, Jaakko Astola, Benoit Vozel, Kacem Chehdi, Marco Carli, Federica Battisti, and C.-C. Jay Kuo. Image database tid2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication*, 30:57 – 77, 2015.
- [18] I. Ram, M. Elad, and I. Cohen. Generalized tree-based wavelet transform. *IEEE Transactions on Signal Processing*, 59(9):4199–4209, Sept 2011.
- [19] R. Reisenhofer, S. Bosse, G. Kutyniok, and T. Wiegand. A Haar Wavelet-Based Perceptual Similarity Index for Image Quality Assessment. *ArXiv e-prints*, July 2016.
- [20] Ron Rubinstein, Michael Zibulevsky, and Michael Elad. Double sparsity: Learning sparse dictionaries for sparse signal approximation. *IEEE Transactions on signal processing*, 58(3):1553–1564, 2010.
- [21] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [22] Lin Zhang, Ying Shen, and Hongyu Li. Vsi: A visual saliency-induced index for perceptual image quality assessment. *IEEE Transactions on Image Processing*, 23(13), 2014.